



ERDC MSRC/PET TR/00-19

**HPC Training Courses at ERDC MSRC
Provided by the Ohio Supercomputer Center**

by

Troy Baer
David Ennis
Leslie Southern

18 May 2000

HPC Training Courses at ERDC MSRC Provided by the Ohio Supercomputer Center

Troy Baer, David Ennis, and Leslie Southern*

Ohio Supercomputer Center (OSC)

1224 Kinnear Road

Columbus, Ohio 43212-1163

** Author for correspondence. E-mail leslie@osc.edu*

Abstract

OSC delivered four high performance computing (HPC) training courses for Engineer Research and Development Center (ERDC) Major Shared Resource Center (MSRC) users during Year 4. These included Using Parallel Libraries; Using the CRAY T3E for Code Development and Analysis; Using the SGI Origin 2000 for Code Development and Analysis; and Parallel Programming using the OpenMP. All courses were delivered at the ERDC MSRC, using TANGO for distance training if possible. This report summarizes the experiences from teaching these classes.

Table of Contents

- I. Introduction and Background
 - II. Course Materials
 - III. Accomplishments
 - IV. Lessons Learned
 - V. Future Directions
 - VI. Acknowledgements
-

Introduction and Background

One of the critical missions of the DoD HPC Modernization Programming Environment and Training (PET) program is to provide user training to assist with the optimum use of scalable parallel computers and related technologies. The CRAY T3E and SGI Origin 2000 are two of the large mainframes resident at ERDC MSRC. It is important that MSRC users be familiar with the hardware as well as with other emerging tools. So, in addition to familiarization courses on these mainframes, OSC taught courses on the use of parallel libraries and OpenMP.

Four courses were delivered to ERDC. The first, on May 11 – 12, 1999, was "Using the CRAY T3E for Code Development and Analysis", by Troy Baer. The second, on

September 8-9, 1999, was “Using Parallel Libraries,” taught by Dr. David Ennis. The third and fourth, on February 28 – March 1, 2000, “Using the SGI Origin 2000 for Code Development and Analysis” and “Parallel Programming using the OpenMP,” were taught by Dr. David Ennis. These courses were taught in the ERDC MSRC training facility, and TANGO Interactive was used to broadcast them to distant attendees.

Course Materials

The CRAY T3E course covered hardware, operating system, operating and programming environments, single-processor optimization, and MPI performance issues. Many of the side discussions concentrated on MPI-related peculiarities. It covered a range of hardware and software features available to users of the CRAY T3E. The topics covered include interactive usage, process management, environment management, batch processing, compilers, debugging tools, performance analysis tools, timing routines, single-node numerical libraries, parallel numerical libraries, fine-grained parallelism, data locality, and streams.

The course entitled “Using Parallel Libraries” was intended to teach supercomputer users about the various libraries available to perform parallel linear algebra calculations without any knowledge of parallel programming languages. Thus it was designed for both the beginning and advanced parallel programmer. The libraries emphasized were PBLACS for basic linear algebra calculations and ScaLAPACK for advanced linear algebra techniques. The course was designed to have a lecture and laboratory component, the latter being critical since an individual really learns programming by trying it, making mistakes, debugging, and finally seeing a solution from the instructor or himself. This material was presented along with the necessary and critical steps that the programmer has to perform in preparation to calling a BLACS or ScaLAPACK routine. In addition, three detailed, complete programs were presented, each one increasing with difficulty.

The Origin 2000 course covered in detail all the aspects of the machine itself: hardware (processor, memory hierarchy, and interconnection topology); operating environment (IRIX operating system, compilers, batch processing); and serial and parallel versions of software tools (debuggers, hardware counters access, and performance analysis).

The OpenMP course described the components (compiler directives and function calls) of this very popular library for loop-level, incremental parallel programming. This course was focused on the OpenMP library itself and its use on multiple MPP platforms.

Accomplishments

The T3E course was broadcast over the DREN to Phillips Laboratory using the TANGO Interactive collaboratory environment software from the Northwest Parallel Architectures Center (NPAC) at Syracuse University. Attendees at Phillips followed along with both the presentation material and an audio/video stream of the instructor's lecture in real time. This was the third synchronous course provided by OSC for the ERDC MSRC and was

technically the smoothest of all. During the 12 hours of training, there was one 5-minute period near the end of the first day when there was a drop in the audio/video conferencing feed.

It was the intent to provide the parallel libraries course to remote users at the Army Research Laboratory (ARL) MSRC via TANGO Interactive. In preparation, course materials for the lectures, and problem sets with solutions for programming sessions were sent ahead of time to both MSRCs. Due to technical problems outside the instructor's control, the remote students were unable to hear the instructor via TANGO. The problems and the attempts to fix them caused the remote students to withdraw from the class after the first day and delayed instruction. The ten remote students represented approximately half the attendance. The problems and the attempts to fix them also created delays in the instruction. All class materials were covered, but no time remained for laboratory programming exercises. The overall evaluation from the students was a 5.0, the highest possible score.

Dr. Ennis is a veteran teacher in the PET program, having taught dozens of courses at various HPC centers. In particular, he has taught the Origin 2000 and OpenMP courses before at ERDC, but combined as a single Origin 2000 workshop. Several changes were made since the last time he taught the course, all of which were significant improvements. The net result was two workshops that ran quite smoothly with no significant problems. The changes were as follows:

- 1) As mentioned, split the existing course in two separate courses. With this approach, students just interested in OpenMP could attend only the last day.
- 2) Movement of the TANGO software from an SGI workstation to an NT workstation. Previously experienced hang-ups, delays, and glitches did not occur during these workshops. Private communication with Mason Colbert at ASC confirmed they were hearing and seeing the lecture and its material all three days.
- 3) An ergonomic improvement made by ERDC was to move the instructor's workstation and teaching location so that it faced the attendees physically present at ERDC (as opposed to looking at them from a ninety-degree angle). When Dr. Ennis addressed students at ASC, he would look into the camera to try making remote eye contact.
- 4) Dr. Ennis scaled down the length of material based on his experience that using TANGO takes up extra lecturing time. One must read a remote question in a chat room, repeat it so that all sites can hear it, and then answer it. In addition, since a shared pointer was not used, Dr. Ennis had to describe the location of significant text on the slide as opposed to just "pointing to it" with a mouse.

Lessons Learned

Thanks to the assistance of many ERDC MSRC PET staff, there was excellent turnaround time for acquiring an MSRC account and SecurID card for the T3E course. During the 12 hours of training, there was one 5-minute period near the end of the first day when there was a drop in the audio/video conferencing feed. This is the third synchronous course provided by OSC for ERDC and was technically the smoothest of all.

As mentioned above, there were problems for TANGO during the parallel libraries course. TANGO continues to improve, but it still occasionally gives us problems for distance training.

The only major problem – and this problem can be insurmountable for course preparation – for the Origin and OpenMP courses was that Dr. Ennis could not get an account on ERDC's Origin 2000 until the morning of the first day of the workshop. Each machine has different characteristics, both at the hardware level and especially at the software level. Dr. Ennis applied for a security clearance by the deadline date, but there is a three-year backup in granting them. This problem needs to be addressed now, and temporary instructor accounts need to be set up months before the actual workshop dates.

Future Directions

For the T3E course, OSC plans to continue improvement to course materials based on attendee questions and comments, and delivery methods. For parallel libraries, similar improvements will be made. The steps mentioned above for the Origin and OpenMP course are examples of how OSC is working with the TANGO environment to improve course delivery.

Acknowledgements

The authors wish to thank the PET and support staff at ERDC MSRC and other recipient sites for their important help and cooperation.